



NRL/MR/5550--01-8591

Mechanisms for Differentiating End-to-End Loss Due to Channel Corruption and Network Congestion

GREGORY W. ROMANIAK
DAVID R. BEERING

*Infinite Global Infrastructures, LLC
Chicago, IL*

MICHAEL A. RUPAR

*Transmission Technology Branch
Information Technology Division*

December 18, 2001

Approved for public release; distribution is unlimited.

20020110 152

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE December 18, 2001		3. REPORT TYPE AND DATES COVERED Interim (January-July 2000)	
4. TITLE AND SUBTITLE Mechanisms for Differentiating End-to-End Loss Due to Channel Corruption and Network Congestion				5. FUNDING NUMBERS PE - 0602232N PR - 01PR02240-00 WU - 55-7384-A-1	
6. AUTHOR(S) Gregory W. Romaniak,* David R. Beering,* and Michael A. Rutar					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory 4555 Overlook Avenue, SW Washington, DC 20375-5320				8. PERFORMING ORGANIZATION REPORT NUMBER NRL/MR/5550--01-8591	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research 800 North Quincy Street Arlington, VA 22217-5660 NASA Glenn Research Center Cleveland, OH				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES *Infinite Global Infrastructures, LLC Chicago, IL					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) <p>This report describes a joint experimental effort by the U.S. Naval Research Laboratory (NRL) and the NASA Glenn Research Center to evaluate standard protocol performance over satellite communication links. The experiment team established a test configuration at the NRL to better understand available mechanisms for the intelligent notification of loss due to channel errors as opposed to loss due to congestion. Several experimental protocol projects that were evaluating error and congestion performance were invited to participate in live satellite tests in the newly created test facility. The experiments discussed ran from March to June of 2000, using the Advanced Communications Technology Satellite (ACTS) and a pair of ACTS Ultra Small Aperture Ka-Band Terminals. During this period, tests were conducted by New Mexico State University on the Space Communications Protocol Standard (SCPS) and by Georgia Tech on a TCP variant known as TCP Peach. The Principal Investigators from those two projects were given access to the NRL test facility via remote access and conducted their protocol experiments themselves. Results of those tests are presented herein.</p> <p>The ACTS spacecraft was decommissioned in May of 2000 and the experiment has been transitioned to Loral Skynet's <i>Telstar II</i> Ku-Band spacecraft, where it continues to the present day. Current areas of investigation are presented at the conclusion of this document.</p>					
14. SUBJECT TERMS Satellite - based network TCP/IP Latency Standard protocol performance				15. NUMBER OF PAGES 17	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL		

CONTENTS

INTRODUCTION	1
BACKGROUND	1
TEST CONFIGURATION	2
EXPERIMENTS PROGRAM	7
CONTINUED WORK	11
SUMMARY	12
REFERENCES	13

MECHANISMS FOR DIFFERENTIATING END-TO-END LOSS DUE TO CHANNEL CORRUPTION AND NETWORK CONGESTION

INTRODUCTION

One of the more challenging aspects of high-performance satellite networking is that TCP/IP, the primary protocol of the modern Internet, is not presently able to distinguish between errors and congestion. Standards-compliant TCP/IP implementations assume that any lost packet is the result of network congestion. TCP/IP employs several algorithms to preserve network stability and to prevent congestion collapse [1]. In terrestrial networks, which typically exhibit very low bit-error rates nearing error-free, this assumption is a valid one. However, in satellite configurations, physical link limitations produce raw link errors, resulting in the loss or damage of the data in transit, requiring retransmissions associated with this loss. Compounding the problem, TCP flow control algorithms cause a halving of the data rate as a result of every detected packet loss. The bursty nature of satellite link errors thus results in a very rapid decline in transmit rate followed by a slow increase that was designed to prevent saturating a congested network. For these satellite data transfer scenarios – especially those at high data rates – this effect can lead to very low link utilization.

One way to mitigate this problem is by improving the satellite channel characteristics. The use of forward error correction coding (FEC) on a satellite link can be used to reduce the Bit-Error Rate (BER). Improving BER through FEC is not always possible, and does not come without cost. FEC requires additional hardware and uses a larger portion of the available channel bandwidth. It can also add delay and timing jitter due to processing overhead [2]. Physical layer methods, including higher gain antennas and amplifiers, can also improve channel characteristics – but the desire for ever-smaller, lower-cost satellite terminals often conflicts with these methods.

These methods, while beneficial, do not address a major TCP shortcoming – its lack of a mechanism to differentiate link errors from congestion. While several proposals exist that address this issue, none has found a consensus recommendation toward the adoption of a standard. The endeavor of this program is twofold – to provide a facility for evaluating experimental protocol enhancements using a real satellite link with deterministic loss performance and to evaluate those enhancements in order to make recommendations on available solutions to this problem to the Internet Engineering Task Force (IETF) standards body.

BACKGROUND

The Satellite and Wireless Networking Section of the Information Technology Division (ITD) at NRL and NASA's Glenn Research Center (GRC) have long been involved in improving the networked connectivity of physical links, both terrestrial and satellite based. GRC was the leader of a consortium of civilian and government laboratories, computing, networking, and commercial satellite industry partners in a multi-year effort to examine high data rate connectivity over satellite and the capabilities of commercial hardware, protocols and applications to support those connections. NRL has for years been developing high data rate satellite and wireless-based networks to support both basic

research and to augment the Navy's communication infrastructure. GRC joined NRL to demonstrate record-setting high data rate connectivity to a ship at sea in 1998, utilizing the ACTS Ka-band satellite to support a 45 Mbps full-duplex networked connection to a vessel on Lake Michigan. This very successful effort underlined the necessity to migrate the commercial standards for IP-based routing in a direction to support links over satellite, which are characterized by periodic link outages and have minimal congestion.

NRL ITD and NASA GRC began a team effort to study high-performance TCP/IP interoperability, aiming to address the errors vs. congestion issue. Necessary for this goal was to establish a stable satellite communications testbed at NRL with predictable performance that could be exercised remotely by experimenters not located at NRL. The team solicited participation from the Internet community through personal invitations and a call for participation distributed via the Internet Engineering Task Force "tcpsat", "pilc", and "end2end" interest groups to participate in this effort. Interest in this work was expressed by several academic institutions, in addition to a number of commercial vendors. Many of the partners involved in the previous work at GRC also offered their support.

The ACTS-based experiment program imposed time and resource constraints on the project, such that this effort was constrained to end at the conclusion of the life of the NASA ACTS program, which concluded in June of 2000. Therefore, only two of the proposals received were invited for immediate participation. Once candidate experimenters were identified, construction, configuration, testing, and characterization of the proposed facility were started at NRL in Washington, DC using NASA-provided Ka-Band Ultra Small Aperture Terminals (USATs). Details of this facility will be presented later in this report. Experimenters were provided with remote access to the computer facilities in Washington, DC via encrypted connections over the Internet, and both experiment teams were able to run a series of experiments and collect useful data prior to the shutdown of the ACTS satellite program [3].

TEST CONFIGURATION

The end-to-end satellite link was characterized to correlate the average bit error rate (BER) performance of the satellite channel to the ratio of received energy per bit to noise, or E_b/N_0 , of the receive-side modem. Through characterization of the satellite link, a table was constructed relating a given E_b/N_0 to its corresponding BER. These data are presented in Figure 1, which shows both laboratory simulation measurements and actual measurements of performance over the ACTS satellite under clear conditions. By varying the output power of a transmit station, the E_b/N_0 at the receive side can be controlled. Therefore, by controlling the E_b/N_0 of the link, we can likewise control the BER. Continuous monitoring of the modem E_b/N_0 allowed adjustment of power levels to account for changes in path conditions (e.g., clouds, rain).




Figure 1: Performance curves (simulated and measured) for the Errors-vs. Congestion TCP/IP satellite testbed at NRL in Washington, DC. The two experimental curves are for a path using the ACTS satellite and two 1.2m Ka-band satellite terminals.

Experiments have used the testbed in a symmetric configuration, to operate with the same data rate and BER in each direction. The satellite link could also be operated in a manner in which BER or data rates are different in each direction. This type of asymmetric configuration could be used to simulate a high-powered, large aperture ground station communicating with a lower-powered, smaller mobile or transportable terminal. In addition, one direction of traffic can be routed over the local terrestrial network, simulating current satellite Internet offerings in which traffic passes over satellite in only one direction, with a terrestrial return link.

Controlling the transmit power (and attenuation at the receive modems) provides a wide range of channel conditions, but it also poses challenges. Satellite links are usually designed to have a certain margin of power to protect against various attenuations in signal levels. Such attenuation can be due to antenna pointing error, atmospheric attenuation due to rain or cloudy conditions, or from variations in ground system performance. This "extra" power available in the end-to-end link is generally referred to as a link margin. The operation of this testbed deliberately removed that margin in order to study behavior near the threshold of acquisition, resulting in desired low E_b/N_0 and thus high BER conditions. Link parameters needed to be continuously tuned to compensate for dynamic channel conditions that caused additional degradation of the link. This is especially the case when operating in the Ka-Band frequency range, as those signals are more drastically affected by weather, especially rain.

To accommodate potentially large fades, the ACTS satellite terminals were designed for a very large link margin, greater than the transmit dynamic range of the modems in the testbed (20 dB), from -25 dBm to -5 dBm. The satellite link was found to produce generally error-free channels, even when operating the modems at their lowest output power settings, due to the link margin designed into the satellite terminals. This required that attenuators be inserted in the cabling between the modems and RF components in order to align the dynamic range of the modems with the range of output power levels required to achieve the desired link characteristics. The attenuation level had to be manually selected, because E_b/N_0 is dependent on the data rate selected for the link. A future upgrade will include programmable attenuators to automate this process.

An important consideration in the general design of the testbed is the issue of remote control. Since the principal investigators in these experiments were distributed geographically, it was imperative that all operations of the terminals could be controlled remotely. The ACTS terminal was designed such that the satellite operators in Cleveland, OH were able to use a modem to dial into the remote terminals to control parameters relating to satellite tracking and RF subsystem controls. Likewise, all computing and networking equipment could be controlled over the Internet via either a direct connection to the laboratory Ethernet or a serial (RS-232) connection to one of the laboratory computing systems. This made it possible for the experiment coordinators not located at NRL, typically in Houston, Cleveland, or Chicago, the satellite operators in Cleveland, as well as the actual experimenters in Georgia or New Mexico to perform experiments on satellite terminals located in Washington, DC.

TCP Error-vs.-Congestion Lab Configuration

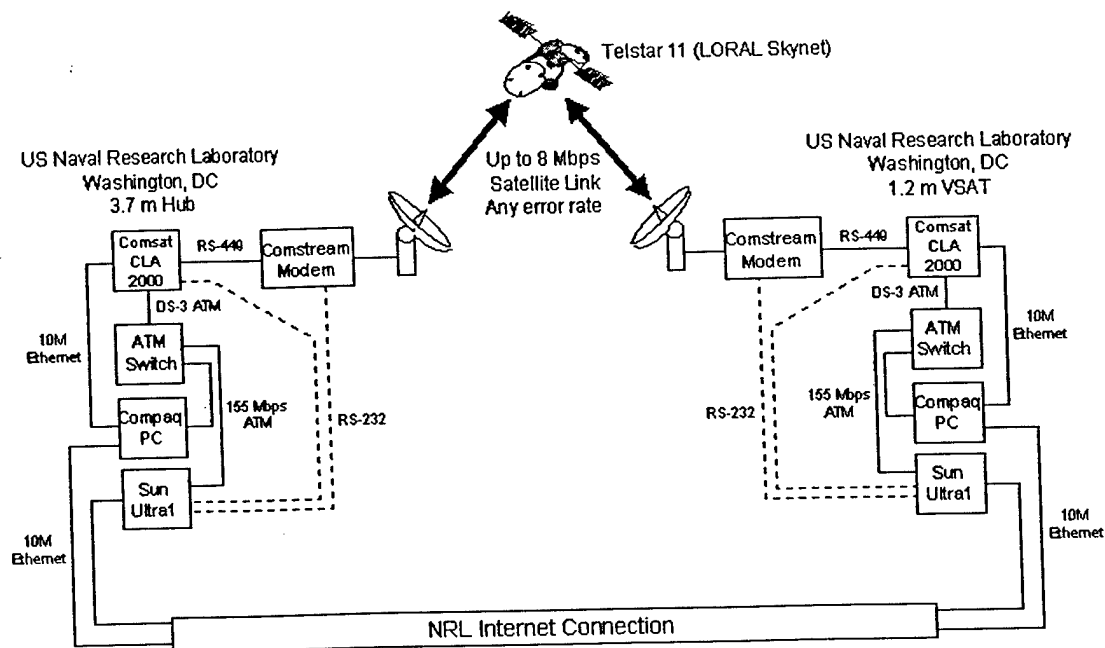


Figure 2. Current test configuration at the US Naval Research Laboratory

The ACTS RF configuration consisted of two NASA-provided Ka-Band Ultra Small Aperture Terminals (USATs) with 1.2 meter diameter antenna systems. Both of these terminals were located on top of the Information Technology Division Building (Bldg. 1) at NRL. One terminal was outfitted with a 1-watt HPA and the other had a newer 2 W unit. This impacted the use of power control since the asymmetric link power required different power settings on each modem to result in the same BER.

By 2000 the ACTS satellite was in an inclined orbit, that is, the location of the satellite was no longer strictly geo-stationary, but moved north and south of its station throughout the day. This required that the pedestals for these antennas were tracking terminals. The tracking system on these terminals used occasional stepping rather than a continuous, smooth movement, causing a small, but noticeable stepping effect on receive signal levels. These variations were generally small enough to be negligible in the experiments.

After the ACTS spacecraft was decommissioned in May of 2000, the team re-configured the experiment to operate on a spacecraft provided by Loral-Skynet – Telstar 11. Access to Telstar 11 was provided as part of an agreement between NRL and Loral-Skynet. The equipment configuration for the Telstar 11 space segment was identical to the configuration used on ACTS. It is worthy to note that since the two antennas presently supporting the experiment at NRL are of substantially different apertures (3.7 meters

versus 1.2 meters) that the same situation still exists regarding power balancing, since the larger aperture has substantially more transmit and receive gain than the smaller aperture.

The modems utilized were Radyne-Comstream - model CM-701s. These were configured with RS-449 high-speed serial terrestrial interfaces and 70 MHz IF RF outputs. The modems are capable of data rates up to 8 Mbps in increments of 1 bps, and were generally configured for 4 Mbps operation in both directions. As mentioned earlier, they also allow for output power control across a 20 dB dynamic range. The modems were configured with Reed-Solomon error correction disabled.

The test network configuration also included a pair of COMSAT Laboratories CLA-2000 units. The devices' primary purpose in the testbed was to perform a conversion of interfaces and data rates between the network-attached devices and the satellite modems. Specific functions supported by the CLA-2000s in this testbed are as follows.

1. Interface conversion from RS-449 (serial interface to the satellite modem) to Ethernet (at 10 Mbps) or ATM (at 45 Mbps). The serial interface operates at data rates up to eight Mbps, and supports asymmetric data rates. The network can therefore be either ATM or Ethernet. The NRL configuration featured devices with both types of network connections.
2. The CLAs can be controlled through the ATM link, the Ethernet link, or the RS-232 serial port. This provides a very flexible network management/telemetry capability for remote configuration and operation.

While the Ethernet port of each CLA-2000 was connected directly to the Ethernet interface of a workstation, the ATM interfaces were connected to a single Marconi ASX-200BX ATM switch via DS-3 coaxial (RG-59) interfaces. The ATM switch, in turn, was connected via OC-3c multimode fiber to each of the four computers in the testbed (shown in Figure 2). The switch allowed the experiment coordinators to isolate each link from the other without requiring two switches by using Permanent Virtual Circuits (PVCs). In addition, it provided the ability to perform terrestrial loop-through testing of initial configurations when the satellite link was unavailable.

The testbed was outfitted with two different computing platforms. Sun Ultra 1 computers, equipped with Marconi SBA-200E ATM interfaces, ran Solaris 7 and supported testing only over ATM because additional Sbus Ethernet interfaces were unavailable. In the future, the Sun hosts will be equipped with additional Ethernet interfaces to support testing of those interfaces across the satellite link. PC platforms were provided on Compaq Proliant hardware with Marconi PCA-200e ATM cards as well as two Ethernet interfaces per computer. This configuration allowed testing with the PC platforms over either ATM or Ethernet while maintaining an out-of-band Ethernet for remote control. The PC hardware platform was initially configured running Redhat Linux with the latest available 2.3.4-*pre* kernel.

EXPERIMENTS PROGRAM

The initial set of tests included two proposed methods to address errored links and an initial baseline of TCP/IP over the satellite channel. The first method is a new algorithm developed at Georgia Institute of Technology called *TCP-Peach*. The second, the Space Communications Protocol Specification (SCPS), is a full protocol stack designed by the Consultative Committee for Space Data Systems (CCSDS) and implemented by Mitre Corporation. Researchers from New Mexico State University who have been testing SCPS performance in their own simulations supported this live satellite testing. Only a subset of the full stack was evaluated, SCPS-TP, the transport protocol mechanism used by SCPS, which is equivalent in function to TCP.

Initial TCP testing was performed in order to develop a baseline upon which to compare the alternative proposals, using the standard TCP stack in Solaris 7. As mentioned earlier, TCP assumes that all packet loss is due to network congestion. As errors occur, TCP reduces the transmit rate and subsequently increases the transmit rate again (slowly) as the number of correctly received packets increases. All TCP baselines are provided on the 154 Experiment web site, located at <http://mrpink.grc.nasa.gov/154>.

Dr. Ian Akyildiz and Dr. Giacomo Morabito, the authors of the algorithm, performed the TCP-Peach [4] testing. TCP-Peach provides a new flow control scheme, the goal of which is to improve throughput performance in satellite networks. It does this through the introduction of three new algorithms called Sudden Start, Rapid Recovery, and Over-Transmit, as well as a modification to the standard TCP Congestion Avoidance algorithm. All the new algorithms are based on the use of low priority "dummy segments". These segments, which carry no new information to the receiver, are used to probe the availability of network resources. Since a congested network that supports the required priority levels will discard these dummy segments, one can infer that the receipt of these segments indicates a lack of congestion. The new algorithms provide a method to more quickly recover from link errors without adversely affecting standard congestion control mechanisms. While the proposed algorithms would eventually be integrated into the operating system's TCP/IP stack, an interim implementation has been developed as a user program running over UDP to demonstrate feasibility. This implementation allows for effective testing of the throughput and the algorithm's response to errors, but cannot yet interact with other TCP streams to perform experimental analysis of congestion control. Simulations of the UDP implementation have shown very good reaction to congestion.

The TCP-Peach testing was conducted on the Sun platforms under Solaris 7, using the satellite link in one direction with an Ethernet connection for the return channel. Experiments were performed with the satellite link configured for Bit Error Rates (BER) of approximately 1×10^{-4} and error-free. Periods of errors during the error-free configuration were due to passing thunderstorms and the associated link attenuation. (While this situation presented itself coincidentally with scheduled testing, a capability to purposely provide this type of channel event is under investigation.) Figures 3 and 4

show histograms of throughput of a single large file in packets/second for $BER=1 \times 10^{-4}$ and an un-degraded link, respectively. Figure 4 demonstrates the effects of a short-duration loss event, that is, a much higher variance in throughput.




Figure 3: A histogram of network performance using TCP Peach over a satellite link, with an operating bit error rate of 1×10^{-4} .




Figure 4, A histogram of network performance using TCP Peach over a "error free" satellite link, with occasional weather events.

These results are very closely correlated with prior simulation results and show an improvement of 96.66% and 92.22% in average throughput over TCP-Reno. A comparison of TCP Peach and TCP Reno (standard) error recovery is shown in Figure 5.

The performance of TCP Peach is the solid line in each plot, and the performance of TCP Reno is the dashed line in each plot. The upper plot depicts the behavior of the congestion window in response to the loss, and the lower plot depicts the number of acknowledged (valid) data segments. The figures depict error recovery performance over a link with a 550 ms roundtrip time (RTT). In each case (TCP Peach and TCP Reno), the congestion window was 30 segments and the experimenters forced the 100th segment to be lost.

The upper plot illustrates that in the case of TCP Peach, the Rapid Recovery mechanism inserts dummy segments that are all acknowledged by the receiver. Immediately following the receipt of the last acknowledgement of a dummy segment, TCP Peach transmits twice the congestion window of data segments. Once those data segments have all been acknowledged, TCP Peach returns the congestion window to the value it had prior to the loss (in approximately 3 RTTs). In this case, TCP Peach is able to make more efficient use of the link.

The lower plot depicts the loss event at the 100th segment. Subsequent to the loss, TCP Peach sends its dummy segments and waits for them to be acknowledged. After the dummy segments and the (2 x cwnd) segments have all been acknowledged, TCP Peach returns to the original configuration of cwnd, outperforming TCP Reno, which is still performing Congestion Avoidance.

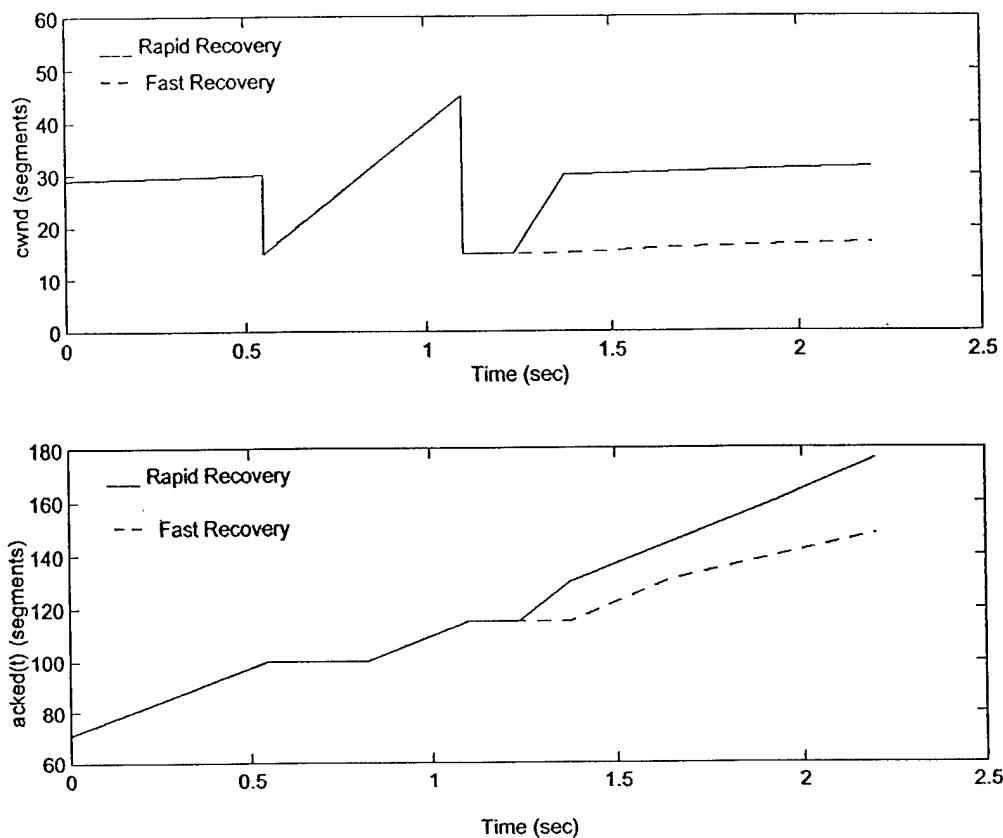


Figure 5: Comparison of error-recovery algorithms in TCP Peach (Rapid Recovery) and TCP-Reno (Fast Recovery).

The second protocol enhancement tested was the Space Communications Protocol Specification (SCPS) [5]. Work on SCPS, developed by the Consultative Committee for Space Data Systems (CCSDS), began in the fall of 1991 and progressed through planned deployment in 1998. It was initially designed for communications in which a spacecraft was one endpoint of a transmission, but the scope was expanded to deal with other scenarios, including networked and non-networked satellite communications environments, cross-linked and non-cross-linked TT&C, and military tactical communications environments. The subset of the SCPS stack tested over the satellite channel was SCPS-TP, the transport protocol of the SCPS suite, which is based on TCP with changes and additions in the following areas:

- Sequence number selection
- Security
- Precedence
- Record Boundary Options
- Best Effort Transport Service (BETS)
- Optional Congestion Control and Corruption Notification
- Detection of Link Outages
- Selective Negative Acknowledgement (SNACK)
- Header Compression

- Multiple Transmission for Forward Error Correction (MFX)

The SCPS reference implementation developed by Mitre Corporation was used in this testing. Several problems were encountered in initial testing due to the limited distribution of the SCPS code. Some of these problems, such as the inability of the implementation to work over ATM interfaces, were never resolved but were worked around. Other problems, namely an addressing problem involving spanning multiple IP subnets and incompatibility with some testing tools were addressed with the help of Mitre. It was also discovered that while the SCPS specification addresses issues of link error and outage mitigation, those features depend on link state information from other layers that was unavailable in this test configuration. Therefore, the supplied SCPS-TP implementation was configured to perform TCP-like congestion control. Further testing will be performed on a specially compiled version of SCPS that assume all packet losses are due to errors.

Testing was performed between the two Compaq PC servers in the testbed by Dr. Stephen Horan and Ru-Hai Wang of New Mexico State University. Initial tests were planned using the ATM interfaces, but after unsuccessful attempts, the system was reconfigured to use Ethernet as the test interface. The test plan consisted of timed file transfers comparing FTP using TCP/IP and SCPS-FP. Repeated file transfers of 1KB, 10KB, 100KB, 1MB, and 10MB files were used to gauge performance at three link states: error free, $BER=1 \times 10^{-5}$, and $BER=1 \times 10^{-6}$. TCPDUMP and TCPTRACE tools were used to record and analyze experiment data. Raw data and additional analysis in the form of a New Mexico State University Technical Report will be available directly from NMSU [6].

CONTINUED WORK

As was stated previously, the testbed has now been re-configured to operate with a Ku-Band channel using Loral Skynet's Telstar 11 spacecraft. At the time of this writing, the Ku-Band configuration using Telstar 11 is still operational. The satellite antennas used for the link are located on the roof of Building 1 at NRL, and are of different apertures (3.7 meters and 1.2 meters), which allows for some interesting additional configurations that more closely model a commercial satellite space segment. The current configuration uses all commercially available satellite equipment as well as the original networking hardware. Both configurations follow the basic network diagram shown previously in Figure 2.

While two proposed methods to differentiate errors from congestion and react accordingly were tested, neither can be recommended as a final solution to the problem. NASA, NRL and their government, academic, and industry partners will continue to seek a solution through continued research into this arena. With the current and future availability of a space segment and human and hardware resources, additional testing is planned and additional participation is being solicited. Current plans include continued

testing of the mechanisms already presented as well as the addition of other identified alternatives.

One such alternative mechanism is Explicit Congestion Notification (ECN), a method by which routers and hosts explicitly communicate a state of network congestion rather than relying on packet loss to infer that state [7]. Cisco Systems has implemented this mechanism in an experimental version of their IOS router software. While ECN does not directly address the issue of link errors, the consortium believes it that the explicit notification of congestion is an important step in differentiating errors from congestion, and is therefore a potential contribution to an overall solution. Work with Cisco Systems to evaluate this protocol enhancement is underway.

Other research in the field has been identified for possible interest in physical testbed evaluation, but contact with the authors has not yet been initiated. The Eifel Algorithm proposes an enhancement to TCP's error recovery scheme, which eliminates retransmission ambiguity, thereby solving the problems caused by spurious timeouts and spurious fast retransmits [8]. Satellite Transport Protocol (STP) is a protocol specifically optimized for the satellite environment in place of TCP, based on an existing ATM-based link layer protocol known as SSCOP. It is suggested that STP can be used as the satellite portion of a split TCP connection and as a transport protocol for control and network management traffic within a satellite communications network [9]. In addition to protocol enhancements, Performance Enhancing Proxies (PEPs) are often cited as a possible solution to the unique problems of satellite links. In addition to these identified possibilities, it is expected that the results presented in this report will prompt additional participation from as-of-yet undetermined parties.

While this group's focus is on enhancing TCP/IP for use over satellite channels, similar problems are encountered in terrestrial wireless applications. Active research in this area should be examined to determine applicability to the general case. In particular, E2E-ELN, presented for wireless terrestrial networks, suggests a method for explicit loss notification using TCP acknowledgements, but assumes a reliable method to determine link state [10]. This and similar work may present methods effective for effectively utilizing both terrestrial and satellite links.

SUMMARY

A combined consortium of NASA, NRL, and research members of the IETF from industry and academia began evaluating the possibility of performing TCP error and congestion differentiation over satellite channels testing early in 2000. The results of testing with three known implementations – standard TCP/IP, TCP Peach and SCPS – helped the researchers to develop important insights into the problem space. While both TCP Peach and SCPS show great promise, it is the opinion of the authors that neither protocol will be sufficient to provide the necessary explicit feedback to the application that losses were due to errors or congestion. In fact, short of the addition of an explicit

physical link layer or data link layer connection to the transport protocol, the end solution may very well be a combination of multiple implicit approaches.

Work in this area will continue in 2001, with additional protocol evaluations, in addition to continued research involving TCP Peach and SCPS. This work will be combined with several additional work areas, including the following:

- Evaluation of Performance Enhancing Proxies (PEPs)
- High-Performance TCP/IP Interoperability
- Internet Protocol over Digital Video Broadcast (DVB)
- Adaptive uplink power control
- Adaptive data rate satellite networks
- Reliable IP Multicast
- Air Interface Options
- Interfaces with other wireless networks (802.11, 3G, etc.)

In the future, the testbed will also be equipped with a pair of Cisco routers to allow support of some router-to-host mechanisms that will be studied during 2001. Most notable among those mechanisms is *Explicit Congestion Notification*, or ECN. The routers will also be used to study TCP/IP performance and efficiency using HDLC as the layer 2 protocol.

All of these work items represent areas of concern for the US Navy, NASA, the Dept. of Defense, partner agencies of the US Government, and the commercial satellite industry. It is the goal of this work to positively impact the availability of standards-based mechanisms for supporting a host of different missions involving satellite components and wireless extensions.

REFERENCES

- [1] Allman, M., Paxson, V. and C.Stevens, "*TCP Congestion Control*", RFC 2581, April 1999
- [2] Allman, M. Glover, D., and L. Sanchez, "*Enhancing TCP over Satellite Channels*", RFC2488, January 1999
- [3] Zernic, M. J., Beering, D. R., and Brooks, D. E., "*NASA and Industry Benefits to ACTS High Speed Network Interoperability Experiments*". Proceedings of 6th Ka Band Utilization Conference, May 2000.
- [4] Akyildiz, I, Morabito, G., and S. Palazzo, "*TCP-Peach: A New Flow Control Scheme for Satellite Networks*". Technical Report. BWN-Lab., Georgia Tech. February 2000.

- [5] Consultative Committee for Space Data Systems, "*Space Communications Protocol Specification (SCPS) – Transport Protocol*." CCSDS 714.0-B-1 Blue Book, May 1999.
- [6] TBD, New Mexico State
- [7] Ramakrishnan, K. and S. Floyd, "*A Proposal to add Explicit Congestion Notification (ECN) to IP*", RFC 2481, January, 1999
- [8] Reiner, L., and R. Katz, "*The Eifel Algorithm: Making TCP Robust against Spurious Retransmissions*". Appears in ACM Computer Communications Review, Vol. 30, No. 1, January 2000.
- [9] Henderson, T. and Katz, R., "*Transport Protocols for Internet-Compatible Satellite Networks*". IEEE Journal on Selected Areas of Communications, 1999.
- [10] Balakrishnan, H., Padmanabhan, V., Seshan, S. and R. Katz, "*A Comparison of Mechanisms for Improving TCP Performance over Wireless Links*". Proceedings of ACM SIGCOMM, 1996.